Incentivizing Users of Data Centers Participate in The Demand Response Programs via Time-Varying Monetary Rewards

Yong Zhan*, Du Xu*, Hongfang Yu*, Shui Yu§
*Key Laboratory of Optical Fiber Sensing and Communications, Ministry of Education,
University of Electronic Science and Technology of China, Chengdu, P. R. China
§School of Information Technology, Deakin University, Victoria, 3125, Australia

arXiv:1604.01950v1 [cs.DC] 7 Apr 2016

***Abstract*—Demand response is widely employed by today's data centers to reduce energy consumption in response to the increasing of electricity cost. To incentivize users of data centers participate in the demand response programs, i.e., breaking the "split incentive" hurdle, some prior researches propose market-based mechanisms such as dynamic pricing and static monetary rewards. However, these mechanisms are either intrusive or unfair. In this paper, we use time-varying rewards to incentivize users, who have flexible deadlines and are willing to trading performance degradation for monetary rewards, grant time-shifting of their requests. With a game-theoretic framework, we model the game between a single data center and its users. Further, we extend our design via integrating it with two other emerging practical demand response strategies: server shutdown and local renewable energy generation. With real-world data traces, we show that a DC with our design can effectively shed its peak electricity load and overall electricity cost without reducing its profit, when comparing it with the current practice where no incentive mechanism is established.**

I. Introduction

Several Demand Response (DR) strategies for Data Centers (DCs) emerged recently, such as time-shifting of workloads [1]. They can effectively reduce/defer a DC's electricity load at specified time in response to the increasing of electricity cost [2]. For example, most of today's DCs need to pay both energy and demand charge, where the demand charge is the cost of the DC's peak electricity load during a billing cycle. In this case, the DCs can shed their demand charges via deferring part of the electricity load from peak hour to off-peak hour [3].

Inevitably, most DR strategies, especially the ones who need workload management, will lead to performance degradation [4]. For example, time-shifting of workloads can greatly increase the latency of service [5]. Naturally, users of DCs are not tend to participate in these DR programs if they are not well incentivized [6]. However, today's DCs usually charge their users via flat-rate based pricing or usage based pricing. Namely, the cost of a user is decided based on the total usage of DC resources, regardless of when the DC resources are utilized [7]. In this case, users do not share the same incentive as the DCs to participate in specified DR programs, which is referred to as "split incentive" [8].

In this paper, we aim at designing a market-based mechanism for a single DC to break the "split incentive" hurdle, which holds the following characteristics:

- *Monetary incentive*: It provides users of DC with incentives to participate in specified DR programs.
- *Non-intrusiveness*: The DR programs are fully voluntary and the costs and performances of users, who decide not to participate in the DR programs, are not effected.
- *Deadline awareness*: All users' requests are processed before their deadlines.
- *Fair reward system*: Only if one participates in the DR programs can gain monetary reward and the reward is proportional to the user's contribution to the DR programs.
- *Simplicity and max-cost guarantee*: The pricing policy, which is used to charge users of DC, should be easy to understand. Meanwhile, it can guarantee that a user's cost for the DC services will not exceed the *baseline cost*.

Note, we assume that every user was previously charged via a commonly-applied pricing policy, e.g., usage based pricing [7], and the corresponding cost is referred to as *baseline cost*.

We use time-varying rewards, on top of the commonly applied pricing policy such as flat-rate based pricing or usage based pricing, to incentivize users of DCs grant time-shifting of their requests for DR. With our design, users, who have tight deadlines and/or are not interested in trading performance degradation for monetary rewards, will be charged by the original pricing policy, and their requests will never be deferred. Namely, it is non-intrusive and deadline aware. Meanwhile, we only reward users, whose requests are deferred, and the reward is proportional to the amounts of requests being deferred. Thus, our reward system is fair. Moreover, our pricing policy and reward system can be easily understand and users' maximum costs with our design will never exceed the *baseline costs*.

With our design, the DC needs to make a suitable trade-off between workload management flexibility and monetary reward. Namely, via giving users higher rewards, the users are more likely to allow the DC to time-shift their requests. On the contrary, if the rewards are relatively low, the DC has less flexibility to response to the electricity price signals. The key contributions of this paper are listed as follows:

- We propose a non-intrusive and fair reward system to in-

centivize users with flexible deadlines grant time-shifting of their requests. In this case, the DC has more workload management flexibility to manage its electricity load and thus reduce its electricity cost.
- With a game-theoretic framework, we model the game between the DC and its users. We first study users' dominant strategies for deciding whether to join in the DR programs. Then, DC's optimal choices of the time-varying rewards are deduced. Specifically, the DC aims at minimizing its electricity cost without reducing of its profit and the problem of electricity cost minimization is formulated via a convex program.
- We extend our design via integrating it with each one of the two emerging practical DR strategies: server shut-down [8]–[10] and local renewable energy generation [6], [11], [12]. Therefore, our design can apply to different combinations of DR strategies.

In the rest of the paper, we first outline several prior work about using market-based mechanisms to incentivize users of DCs participate in the DR programs in section II. In section III, user's surplus and DC's profit are formulated. We model the game between users and DC in section IV and extend our design in section V. We evaluate our design with real-world data traces and conclude this paper in section VI and VII, respectively.

II. Related Works

With specified DR programs, DC can effectively reduce its electricity load [4], as well as the demand charge and the overall electricity cost [3], [13]. Considering that most DR strategies involve management of users' workloads, e.g., time-shifting of requests, several market-based mechanisms have been proposed to incentivize users grant workload management [6], [8], [14]–[16].

[6] propagated the overall electricity cost onto the users' costs via charging users by Time-of-Use (ToU) pricing. Namely, at each period of time, the price of an instance of DC, i.e., a combination of DC resources such as CPU, memory and storage, is set to be proportional to the prediction of the DC's electricity load. With this pricing policy, users, who can tolerate performance degradation and are sensitive to price, may reduce their purchases of instances at peak hour in response to the increasing of cost. In this case, the DC's peak electricity load and its demand charge can be reduced.

[14] fairly distributed a DC's overall electricity cost among its users. With [14]'s design, if the DC's electricity cost is mainly decided by its peak electricity load, users of the DC can gain considerable cost reduction via reducing its demand when an overall peak occurs. Thus, users are well incentivized to participate in the DR programs to shed the DC's peak electricity load.

However, both [6] and [14] are intrusive since the costs of all users are effected regardless of whether they are interested in participating in the DR programs. Meanwhile, they do not guarantee that the cost of a user with their designs will never exceed the *baseline cost*.

[8] proposed Usage-based Pricing with Monetary Reward (UPMR). [8]'s mechanism is non-intrusive and can provide max-cost guarantee. At the beginning of each billing cycle, DC releases the reward function to its users. Then, each user can make the decision of whether to join the DR programs in the upcoming billing cycle based on the reward function and its sensitivity to the performance degradation. The key difference between [8]'s work and our design is that: the monetary reward offered by [8] is static, i.e., it will remain unchanged during a billing cycle, and the reward is given to all users, who are willing to participate in the DR programs, regardless of whether their workloads are modified for DR. Thus, [8]'s reward system is not fair. On the contrary, In this paper, we provide users with time-varying monetary rewards and users can change their decisions of whether to join the DR programs during a billing cycle based on the real-time reward. In this case, the amount of users, who have joined in the DR programs, is time-varying, and thus the DC embraces more load flexibility. Moreover, the reward given to each user with our design is proportional to the user's contribution to the DR program, i.e., the amount of requests being deferred, which is more fair.

III. Problem Formulation

A time-slotted system is used in this paper. Namely, we divide a billing cycle into τ time slots, with equally length of T (hour). In this section, we formulate a user's surplus, as well as a DC's profit.

A. User's Surplus

Let $V_i[t]$ denote the net utility can be gained by user i after processing one of its request generated at time slot t without participating in any DR programs, where $i \in \{1, \cdots, I\}$ denotes the index for users and $t \in \{1, \cdots, \tau\}$ denotes the index for time slots. Here, I denotes the number of active users. Let $\delta_i[t]$ (\$) denote the average *baseline cost* for processing one request generated by user i at time slot t. Further, let $\kappa_i[t]$ and $\gamma[t]$ denote the net utility loss and financial reward caused by modifying one of user i's request generated at time slot t, respectively.

We formulate user i's surplus after processing one of its request generated at time slot t, which is denoted by $S_i[t]$, as below:

$$S_i[t] = V_i[t] - \delta_i[t] + \mathbb{1}_i[t](\gamma[t] - \kappa_i[t]), \tag{1}$$

where $\mathbb{1}_i[t]$ is a binary variable and indicates whether the request is modified for DR. Namely, if $\mathbb{1}_i[t] = 1$, the request is modified, and thus user i will loss $\kappa_i[t]$ net utility but enjoy $\gamma[t]$ monetary reward. On the contrary, if $\mathbb{1}_i[t] = 0$, the request has not been modified and thus no net utility is lost and no reward is given.

B. DC's Profit

Revenue: Suppose that the DC's revenue is constituted by all of its active users' *baseline costs*. Thus, we formulate the

DC's revenue over a billing cycle via

$$\text{Revenue} = \sum_{t=1}^{\tau} \sum_{i=1}^{I} \lambda_i[t]\delta_i[t], \tag{2}$$

where $\lambda_i[t]$ denotes the amount of requests generated by user i at time slot t.

Electricity cost: The electricity cost of today's DC is usually constituted by two parts: energy charge and demand charge, where energy charge is decided by the DC's total energy consumption and demand charge is computed based on the DC's peak electricity load. Specifically, the demand charge can weigh heavily on a DC's overall electricity cost [13], [17]. Let $\alpha[t]$ ($/KWh) denote the price of energy charge at time slot t and β_j ($/KW) denote the price of type-j demand charge. Here, $j \in \{1, \cdots, J\}$ is the index for types of demand charge and J measures the total types of demand charge. For example, in [3], the authors divided the whole billing cycle into on-peak and off-peak periods and the prices of demand charge during on-peak and off-peak periods can be different. Therefore, we formulate the DC's electricity cost over a billing cycle via

$$\text{Cost} = \sum_{t=1}^{\tau} T\alpha[t]P[t] + \sum_{j=1}^{J} \max_{t \in A_j} \beta_j P[t], \tag{3}$$

where A_j denotes the periods of time falling into the type-j demand charge window. Here, $P[t]$ (KW) denotes the average electricity load of the DC at time slot t.

Like in [9], the average electricity load of the DC at time slot t is formulated as

$$P[t] = E_{pue}[t] \sum_{n=1}^{N} (e_0 + u_n[t]e_1), \tag{4}$$

where $E_{pue}[t]$ denotes the average ratio of overall DC energy consumption to IT energy usage at time slot t, $n \in \{1, \cdots, N\}$ is the index for active servers, N denotes the total amount of active servers, e_0 (KW) denotes the power of an idle server, e_1 (KW) measures the increase of power of a server with unit increase in CPU utilization and $u_n[t]$ denotes the average CPU utilization of server n at time slot t.

Without loss of generality, we assume that each request needs the same amount of DC resources to be fully processed and each server can handle the same amount of requests within a time slot. Meanwhile, according to [18], the optimal workload dispatching rule is to distribute all requests equally on all active servers. Thus, we formulate the average CPU utilization ratio of server n at time slot t via

$$u_n[t] = \frac{\hat{\lambda}[t]}{N\nu}, \tag{5}$$

where $\hat{\lambda}[t]$ denotes the aggregate requests scheduled at time slot t and ν denotes the maximum amount of requests can be processed by a server within a time slot in fully utilized condition.

We use time-shifting of users' requests in this paper to shed the DC's peak electricity load. Let D denote the maximum length of deferrable periods of time. Namely, all requests generated at time slot t should be scheduled no later than time slot $t + D$. Let $\eta_d[t]$ denote the aggregate requests generated at time slot t and scheduled at time slot $t + d$, where $d \in \{0, \cdots, D\}$ is the index for the amount of time slots of request being deferred. Specifically, $\eta_0[t]$ measures the amount of requests generated at time slot t and scheduled at the same time slot, i.e., they are not deferred. On the contrary, $\sum_{d=1}^{D} \eta_d[t]$ measures the amount of deferred requests generated at time slot t. Thus, we formulate the aggregate requests scheduled at time slot t via

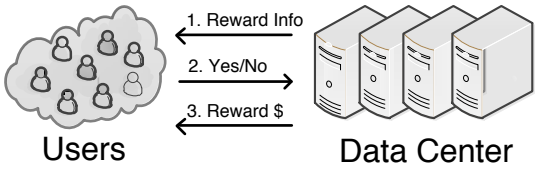

Fig. 1. Outline of the interactions between users and DC.

$$\hat{\lambda}[t] = \sum_{d=0}^{D} \eta_d[t-d]. \tag{6}$$

Reward: We give monetary rewards to users for allowing deferring their requests and the rewards are proportional to the amounts of requests being deferred. Thus, we formulate the total amount of rewards given to the users over a billing cycle via

$$\text{Reward} = \sum_{t=1}^{\tau} \gamma[t] \sum_{d=1}^{D} \eta_d[t], \tag{7}$$

Profit: We define the profit of the DC over a billing cycle via

$$\text{Profit} = \text{Revenue} - \text{Cost} - \text{Reward}. \tag{8}$$

Note, the DR programs have no impact on the Revenue term, since it is fully constituted by the *baseline costs*. Thus, we treat the Revenue term as a given constant in the rest of the paper.

IV. Game-Theoretic Framework

In this paper, we assume that the DC has perfect knowledge of the aggregate requests in the upcoming billing cycle. A game-theoretic framework is proposed for the DC to deduce the optimal choices of rewards in this section. The interactions between users and DC is shown in Fig. 1. At the beginning of each time slot, DC first releases the reward information, i.e., $\gamma[t]$, to users. Next, users answer whether they are willing to allow DC to defer their requests generated at this time slot. Finally, DC schedules users' requests and rewards users whose requests are deferred.

A. Users' Dominant Strategies

Every user needs to make the decision of whether to allow DC to defer its requests generated at the upcoming time slot based on the reward, i.e., $\gamma[t]$, and its sensitivity to time-shifting of its requests, i.e., $\kappa_i[t]$.

Assumption 1: Suppose every user is *rational*, which means that it wants to maximize its surplus as defined in (1).

Assumption 2: Suppose that the DC holds an enormous amount of users. In this case, each user's contribution to the DC's overall electricity load is negligible, and thus every user cannot strategically manipulate the price of DC services and the rewards given to users. Namely, the users are *price takers*.

Theorem 1: From Assumption 1 and 2, $\forall i, t$, if the reward exceeds the net utility loss caused by time-shifting of a user's requests, i.e., $\gamma[t] > \kappa_i[t]$, the dominant strategy for the user is to grant time-shifting of its requests. On the contrary, $\forall i, t$, if $\gamma[t] \leq \kappa_i[t]$, the dominant strategy for the user is not to participate in the DR programs.

The proof of Theorem A can be found in Appendix A.

B. DC's Optimal Decisions

Variables: The DC needs to make two decisions at each time slot: $\gamma[t]$ and $\eta_d[t]$. Specifically, $\gamma[t]$ is about the reward given to users for granting time-shifting of their requests and $\eta_d[t]$ elaborates the detail of request distribution.

Constraints: First, we define the domain of the variables:

$$Lb[t] \leq \gamma[t] \leq Ub[t], \quad \forall t, \tag{9}$$

$$\eta_d[t] \geq 0, \quad \forall d, t, \tag{10}$$

where $Lb[t]$ and $Ub[t]$ denote the lower bound and upper bound of net utility loss factor, i.e., $\kappa_i[t]$, of elastic users, respectively. Here, the elastic users are the ones who have flexible deadlines and are will allow DC to time-shift their requests if they are well incentivized. On the contrary, the users, who have tight deadlines or/and are not interested in trading performance degradation for monetary rewards, are referred to as the inelastic users. Specifically, for all inelastic users, we let their net utility loss factor $\kappa_i[t]$ to be infinity. In this case, $\forall \gamma[t] \in [0, Lb[t]]$, there is no user will join the DR programs. Meanwhile, $\forall \gamma[t] \in [Ub[t], \infty)$, all elastic users are willing to allow DC to time-shift their requests. Thus, in (9), we indicate that the DC has no interest to set the reward lower than $Lb[t]$ or higher than $Ub[t]$.

Second, the amount of requests scheduled at each time slot should not exceed the capacity of the DC, which is formulated as

$$\sum_{d=0}^{D} \eta_d[t-d] \leq N\nu. \tag{11}$$

With (11), we ensure that the average CPU utilization ratio of each server of the DC defined by (5) will never exceed 1.

Third, all requests generated at each time slot should be scheduled before the deadline, which is formulated as

$$\sum_{d=0}^{D} \eta_d[t] = \lambda[t], \quad \forall t, \tag{12}$$

where $\lambda[t]$ denotes the aggregate requests of users generated at time slot t and

$$\lambda[t] = \sum_{i=1}^{I} \lambda_i[t], \quad \forall t. \tag{13}$$

Fourth, we notice that the uniform distribution is widely used in economic models [19]. Thus, we assume that, $\forall t$, the net utility loss factors $\kappa_i[t]$ of elastic users are uniformly distributed from $Lb[t]$ to $Ub[t]$. Moreover, without loss of generality, suppose that each user will generate the same amount of requests at each time slot, i.e., $\forall t$ and $\forall i^1, i^2 \in \{1, \cdots, I\}$, $\lambda_{i^1}[t] = \lambda_{i^2}[t]$. From Theorem 1, we obtain that, $\forall t$, the amount of deferrable requests, i.e., the sum of requests generated by users who grant time-shifting of their requests, equals $\pi[t]\lambda[t]\frac{\gamma[t]-Lb[t]}{Ub[t]-Lb[t]}$, where $\pi[t]$ denotes the ratio of the elastic users to the overall users at time slot t. Obviously, the amount of requests being deferred should be no higher than the sum of deferrable requests, which is modeled as

$$\sum_{d=1}^{D} \eta_d[t] \leq \pi[t]\lambda[t]\frac{\gamma[t]-Lb[t]}{Ub[t]-Lb[t]}, \quad \forall t. \tag{14}$$

Last but not least, in this paper, we aim at minimizing the electricity cost of the DC while ensuring that the profit with our design is no less than the one without any DR programs and incentive mechanisms, which is modeled as

$$\text{Revenue} - \text{Cost} - \text{Reward} \geq \text{Revenue} - \text{Baseline Cost}, \tag{15}$$

where the Revenue, Cost and Reward terms are defined by (2), (3) and (7), respectively. Here, the Baseline Cost term is the electricity cost of the DC if no DR strategy is implemented. Correspondingly, there is also no reward given to the users. Further, we simplify (15) and reformulate it as

$$\text{Cost} + \text{Reward} \leq \text{Baseline Cost}. \tag{16}$$

Minimization of electricity cost: We define the problem of minimization of the DC's electricity cost without reducing its profit via

$$\begin{aligned} \min_{\gamma[t],\eta_d[t]} \quad & \text{Cost} \\ \text{s.t.} \quad & (9)-(12), (14), (16). \end{aligned} \tag{17}$$

where the Cost term is defined in (3).

Apparently, problem (17) is not convex caused by the non-convex term Reward (7).

Theorem 2: For any given $\eta_d[t]$, $\gamma^*[t]$, which is defined as

$$\gamma^*[t] = \frac{(Ub[t]-Lb[t])\sum_{d=1}^{D}\eta_d[t]}{\pi[t]\lambda[t]} + Lb[t], \quad \forall t, \tag{18}$$

is one of the optimal choice of $\gamma[t]$ of problem (17).

The proof of Theorem 2 can be found in Appendix B. From Theorem 2, we let $\gamma[t] = \gamma^*[t]$, $\forall t$. Further, we update the problem (17) to be

$$\begin{aligned} \min_{\eta_d[t]} \quad & \text{Cost} \\ \text{s.t.} \quad & (10)-(12), (16). \end{aligned} \tag{19}$$

Apparently, problem (19) is convex, which can be effectively solved via a convex tool such as CVX [20].

V. Extensions

We extend our design via integrating workload time-shifting with another emerging practical DR strategy: server shutdown [8]–[10] or local renewable energy generation [6], [11], [12] in this section.

A. Server Shutdown

Considering that the proportionality degree of today's commercial server is relatively low, e.g., an idle server can consume over 85% of peak power [21], shutting down the idle servers can save a huge amount of energy. In this paper, we take into consideration the energy overhead and the wear-and-tear cost caused by turning on/off servers as in [8]. In this case, at each time slot, the DC also needs to decide how many servers should be switched on/off, which is denoted by $m_{on}[t]$ and $m_{off}[t]$, respectively.

As in [8], we define the electricity cost of the DC over a billing cycle with server shutdown, which is denoted by Cost_S, as

$$\begin{aligned}\text{Cost}_S = &\sum_{t=1}^{\tau} \alpha[t](TP_s[t] + E_{pue}P_o[t]) + \\ &\sum_{j=1}^{J} \beta_j \max_{t\in A_j}\left(P_s[t] + \frac{E_{pue}P_o[t]}{T}\right),\end{aligned} \tag{20}$$

where $P_s[t]$ and $P_o[t]$ denote the average electricity load of the DC with server shutdown and the average energy overhead of turning on/off servers at time slot t, respectively. Here, both of the definitions of $P_s[t]$ and $P_o[t]$ can be found in [8]. Moreover, the wear-and-tear cost of servers over a billing cycle, which is denoted by Wear, is also defined as in [8].

The problem of minimizing the DC's electricity cost with server shutdown without reducing its profit is modeled as

$$\begin{aligned}\min_{\substack{\eta_d[t],\\ m_{on}[t],\\ m_{off}[t]}} \quad & \text{Cost}_S \\ \text{s.t.} \quad & (10), (12), \\ & \text{Cost}_S + \text{Reward} + \text{Wear} \le \text{Baseline Cost}, \\ & m_{on}[t] \ge 0, \quad m_{off}[t] \ge 0, \qquad \forall t, \\ & \frac{\hat{\lambda}[t]}{\nu} \le m[0] + \sum_{t'=1}^{t}(m_{on}[t] - m_{off}[t]) \le N, \quad \forall t,\end{aligned} \tag{21}$$

where $m[0]$ denotes the amount of switched on servers at the beginning of the billing cycle. Here, the last constraint of problem (21) indicates that the amount of switched on servers at each time slot should be large enough to handle all scheduled requests, but no larger than the quantity of available servers, which is denoted by N. Like problem (19), problem (21) is also convex.

B. Local Renewable Energy Generation

Recently, many state-of-the-art DCs have invested heavily to develop local renewable energy generation, such as Apple and Google [4]. Let $G[t]$ (KWh) denote the renewable energy generated at time slot t. As in [6], we assume that the information of renewable energy can be perfectly predicted by the DC and the part of renewable energy, which is not used by the DC, is dropped. Thus, the electricity cost with local renewable energy generation, which is referred to as Cost_R, is defined by

$$\begin{aligned}\text{Cost}_R = &\sum_{t=1}^{\tau} \alpha[t]\left[TP[t] - G[t]\right]^+ + \\ &\sum_{j=1}^{J} \beta_j \max_{t\in A_j}\left[P[t] - \frac{G[t]}{T}\right]^+,\end{aligned} \tag{22}$$

where $[x]^+$ represents the function $\max\{x, 0\}$.

Next, we model the problem of minimization of the DC's electricity cost with local renewable energy generation without reducing its profit via

$$\begin{aligned}\min_{\eta_d[t]} \quad & \text{Cost}_R \\ \text{s.t.} \quad & (10) - (12), \\ & \text{Cost}_R + \text{Reward} \le \text{Baseline Cost},\end{aligned} \tag{23}$$

which is also convex.

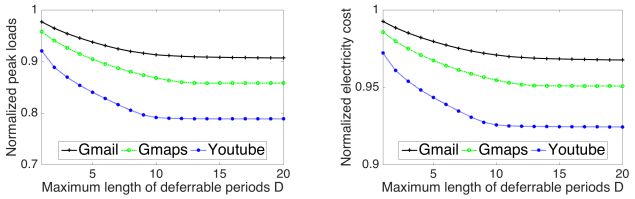

(a) (b)

Fig. 2. Normalized results of a DC with our design: (a) Peak electricity load, (b) Electricity cost.

VI. Case Studies

Three real-world data traces are used in this paper to represent users' requests: Gmail U.S., Gmaps U.S. and Youtube U.S. from January 1, 2014 to January 31, 2014 [22]. According to the electric rates of industrial power service offered by South Carolina Electric and Gas [23], $\forall t$, let $\alpha[t] = \$0.05207$ per KWh, $J = 1$ and $\beta_1 = \$15.59$ per KW. Let $T = 1$ (hour) and $\tau = 720$. Namely, the billing cycle is 30 days basis. Let $e_0 = 0.1$ KW, $e_1 = 0.1$ KW, $\nu = 20$ and $E_{pue} = 1.2$. $\forall t$, let $\pi[t] = 0.5$, $Lb[t] = 10e^{-4}$ (\$) and $Lb[t] = 10e^{-3}$ (\$). Next, in this section, we first evaluate our work by comparing its performance with the one without any incentive mechanisms, which represents the current practice of DC and is referred to as **Baseline**. Then, we study the impact of integrating our design with another DR strategy on the reduction of peak electricity load and overall electricity cost.

A. Performance Evaluation

We study the DC's peak electricity load and electricity cost with our design under different maximum lengths of deferrable

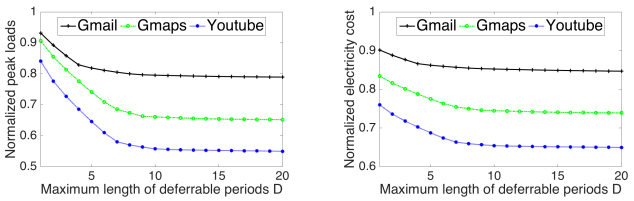

(a) (b)

Fig. 3. Normalized results of a DC with our design and server shutdown: (a) Peak electricity load, (b) Electricity cost.

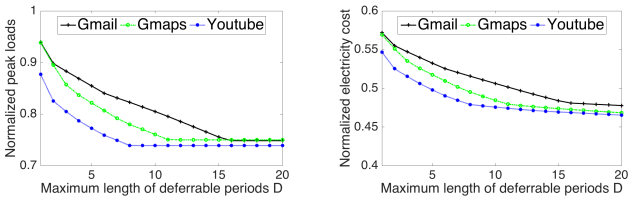

(a) (b)

Fig. 4. Normalized results of a DC with our design and local renewable energy generation: (a) Peak electricity load, (b) Electricity cost.

periods of time, i.e., D. Fig. 2 shows the normalized results with respect to the results of Baseline.

Three information can be found from Fig. 2: 1) With the increasing of the maximum length of deferrable periods of time, our design can shed the DC's peak electricity load and reduce its overall electricity cost more effectively. 2) When the maximum length of deferrable periods is relatively large, e.g., larger than 10 time slots as shown in Fig. 2, the effectiveness of our design cannot be greatly increased via further increasing the maximum length of deferrable periods of time. 3) DC with our design always outperforms the one without any DR programs and incentive mechanisms. For example, when $D = 10$, the DC with data traces of Youtube U.S. can shed its peak electricity load by 20.9% and reduce its overall electricity cost by 7.4%.

B. Impact of Server Shutdown and Local Renewable Energy Generation

We integrate our design with another emerging practical DR strategy: server shutdown or local renewable energy generation, and evaluate their performances. Fig. 3 and Fig. 4 show the normalized results of a DC with server shutdown and local renewable energy generation, respectively. Here, the base for both normalizations is the results of Baseline. Specifically, the DC as shown in Fig. 4 is equipped with several wind turbines and the wind speed data is gathered from January 1, 2012 to January 30, 2012 and is available in [24]. By analyzing Fig. 3 and 4, the three information found in section VI-A still hold. Moreover, By comparing Fig. 3 and 4 against Fig. 2, we find that the peak electricity load and the overall electricity cost of a DC with our design can be further reduced via integrating with other DR strategies such as server shutdown and local renewable energy generation.

VII. Conclusion

In this paper, we proposed time-varying rewards to incentivize users of DCs grant time-shifting of their requests. Our design is non-intrusive, deadline aware and easy to understand. Moreover, our reward system is fair and we can provide users with max-cost guarantee. With a game-theoretic framework, we modeled the game between users and a single DC and formulated/solved the problem of minimization of the DC's electricity cost without reducing its profit. With real-world data traces, we showed that our design can greatly reduce the DC's peak electricity load and overall electricity cost. We also extended our design to another emerging practical scenario where the DC has employed workload time-shifting and server shutdown or local renewable energy generation simultaneously. Accordingly, our design can apply to a DC with different combinations of DR strategies.

Appendix A
Proof of Theorem 1

Apparently, $\forall i, t$, from (1), if a user does not grant time-shifting of its requests, its surplus at time time slot t, which is referred to as $S_i^{No}[t]$, equals

$$S_i^{No}[t] = \lambda_i[t](V_i[t] - \delta_i[t]). \tag{24}$$

On the contrary, if the user decides to participate in the DR programs, its surplus at time slot t, which is referred to as $S_i^{Yes}[t]$, is

$$S_i^{Yes}[t] = \lambda_i[t](V_i[t] - \delta_i[t]) + \hat{\lambda}_i[t](\gamma[t] - \kappa_i[t]), \tag{25}$$

where $\hat{\lambda}_i[t]$ denotes amount of deferred requests generated by user i at time slot t. Note, since DC makes the decision of whether to defer part or all of the users requests if the user grant time-shifting of its requests, the user cannot perfectly predict how many requests will be deferred but can ensure that $0 \leq \hat{\lambda}_i[t] \leq \lambda_i[t]$.

Next, we prove that Theorem 1 is correct in two complementary cases: First, $\forall i, t$, if $\gamma[t] > \kappa_i[t]$, $S_i^{No}[t] \leq S_i^{Yes}[t]$. Namely, in this case, the dominant strategy for the user is to join in the DR programs. Second, $\forall i, t$, if $\gamma[t] \leq \kappa_i[t]$, $S_i^{No}[t] \geq S_i^{Yes}[t]$, and thus the optimal choice for the user is not to participate in the DR program in this case. ■

Appendix B
Proof of Theorem 2

First, we can easily obtain that $\forall t$, $\gamma^*[t] \geq Lb[t]$. Second, from (14), we have $\forall t$,

$$\sum_{d=1}^{D} \eta_d[t] \leq \pi[t]\lambda[t] \quad \Longrightarrow$$

$$\frac{(Ub[t] - Lb[t]) \sum_{d=1}^{D} \eta_d[t]}{\pi[t]\lambda[t]} \leq Ub[t] - Lb[t] \quad \Longrightarrow$$

$$\gamma^*[t] \leq Ub[t].$$

Meanwhile, $\gamma^*[t]$ satisfy constraint (14) of problem (17). Thus, $\gamma^*[t]$ is a feasible solution of $\gamma[t]$ of problem (17). Next,

we prove Theorem 2 by contradiction. Assume that $\gamma^*[t]$ is not an optimal solution of $\gamma[t]$ of problem (17) and let $\gamma^c[t]$ denote the true optimal solution of $\gamma[t]$ of problem (17). From (3), $\text{Cost}|_{\gamma[t]=\gamma^c[t]} = \text{Cost}|_{\gamma[t]=\gamma^*[t]}$, which contradicts the assumption that $\gamma^*[t]$ is not an optimal solution of $\gamma[t]$ of problem (17). ■

REFERENCES

[1] G. Ghatikar, "Demand response opportunities and enabling technologies for data centers: Findings from field studies," tech. rep., PG&E/SDG&E/CEC/LBNL, 2014.

[2] A. Wierman, Z. Liu, I. Liu, and H. Mohsenian-Rad, "Opportunities and challenges for data center demand response," in *Proc. of IEEE IGSC*, Dallas, TX, Nov 2014.

[3] V. Mathew, R. Sitaraman, and P. Shenoy, "Reducing energy costs in internet-scale distributed systems using load shifting," in *Proc. of IEEE COMSNETS*, Bangalore, India, Jan 2014.

[4] Z. Liu, I. Liu, S. Low, and A. Wierman, "Pricing data center demand response," *SIGMETRICS Perform. Eval. Rev.*, vol. 42, pp. 111–123, Jun 2014.

[5] C. Wang, B. Urgaonkar, Q. Wang, and G. Kesidis, "A hierarchical demand response framework for data center power cost optimization under real-world electricity pricing," in *Proc. of IEEE MASCOTS*, Paris, France, Sep 2014.

[6] Y. Zhan, M. Ghamkhari, D. Xu, and H. Mohsenian-Rad, "Propagating electricity bill onto cloud tenants: Using a novel pricing mechanism," in *Proc. of IEEE Globecom*, San Diego, CA, Dec 2015.

[7] A. Li, X. Yang, S. Kandula, and M. Zhang, "Cloudcmp: comparing public cloud providers," in *Proc. of ACM IMC*, Melbourne, Australia, Nov 2010.

[8] Y. Zhan, M. Ghamkhari, D. Xu, S. Ren, and H. Mohsenian-Rad, "Extending demand response to tenants in cloud data centers via non-intrusive workload flexibility pricing," *arXiv:1603.05746*, Mar 2016.

[9] D. Paul, W.-D. Zhong, and S. K. Bose, "Demand response in data centers through energy-efficient scheduling and simple incentivization," *IEEE Systems Journal*, vol. PP, pp. 1 – 12, Sep 2015.

[10] J. Li, Z. Bao, and Z. Li, "Modeling demand response capability by internet data centers processing batch computing jobs," *IEEE Trans. on Smart Grid*, vol. 6, pp. 737–747, Mar 2015.

[11] W. Deng, F. Liu, H. Jin, B. Li, and D. Li, "Harnessing renewable energy in cloud datacenters: opportunities and challenges," *IEEE Network*, vol. 28, pp. 48–55, Jan 2014.

[12] M. Ghamkhari and H. Mohsenian-Rad, "Energy and performance management of green data centers: A profit maximization approach," *IEEE Transactions on Smart Grid*, vol. 4, pp. 1017–1025, Jun 2013.

[13] H. Xu and B. Li, "Reducing electricity demand charge for data centers with partial execution," in *Proc. of ACM e-Energy*, Cambridge, UK, Jun 2014.

[14] N. Nasiriani, C. Wang, G. Kesidis, B. Urgaonkar, L. Y. Chen, and R. Birke, "On fair attribution of costs under peak-based pricing to cloud tenants," in *Proc. of IEEE MASCOTS*, Atlanta, GA, Sep 2015.

[15] C. Wang, N. Nasiriani, G. Kesidis, B. Urgaonkar, Q. Wang, L. Y. Chen, A. Gupta, and R. Birke, "Recouping energy costs from cloud tenants," in *Proc. of ACM e-Energy*, Bangalore, India, Jul 2015.

[16] C. Wang, B. Urgaonkar, G. Kesidis, U. V. Shanbhag, and Q. Wang, "A case for virtualizing the electric utility in cloud data centers," in *Proc. of USENIX HotCloud*, Philadelphia, PA, Jun 2014.

[17] M. A. Islam, H. Mahmud, S. Ren, and X. Wang, "Paying to save: Reducing cost of colocation data center via rewards," in *Proc. of IEEE HPCA*, Burlingame, CA, Feb 2015.

[18] M. Lin, A. Wierman, L. L. Andrew, and E. Thereska, "Dynamic right-sizing for power-proportional data centers," *IEEE/ACM Trans. on Networking*, vol. 21, pp. 1378–1391, Oct 2013.

[19] R. Gibbens, R. Mason, and R. Steinberg, "Internet service classes under competition," *IEEE Journal on Selected Areas in Communications*, vol. 18, pp. 2490–2498, Dec 2000.

[20] "Cvx: Matlab software for disciplined convex programming." http://cvxr.com/cvx/.

[21] A. Qureshi, *Power-demand routing in massive geo-distributed systems*. PhD thesis, Massachusetts Institute of Technology, 2010.

[22] "Browse real-time traffic to google products and services." http://www.google.com/transparencyreport/traffic/explorer.

[23] "SCE&G electric rates." https://www.sceg.com/for-my-business/manage-my-service/rates.

[24] "Wind energy." http://www.windenergy.org.